

What is claimed is:

1. A method for processing a memory instruction, the method comprising:
 - obtaining a memory instruction;
 - obtaining one or more memory address operands;
 - 5 creating a virtual memory address using the one or more memory address operands;
 - translating the virtual memory address into a physical memory address; and
 - executing the memory instruction using a cache controller,
 - wherein the cache controller uses the memory instruction and the physical
- 10 memory address to determine whether to access a portion of a local or a remote cache.

2. The method of claim 1, wherein the obtaining of the memory instruction includes obtaining a scalar memory instruction.
- 15 3. The method of claim 1, wherein the obtaining of the memory instruction includes obtaining a memory instruction selected from a group consisting of a scalar load instruction, a scalar store instruction, a prefetch instruction, a synchronization instruction, and an atomic memory operation (AMO) instruction.
- 20 4. The method of claim 1, wherein the obtaining of the memory instruction includes obtaining a scalar store instruction, and wherein the method further comprises obtaining scalar store data.
- 25 5. The method of claim 1, wherein the translating includes translating the virtual memory address into a physical memory address using a translation look-aside buffer (TLB).

6. The method of claim 1, wherein the method further comprises committing the memory instruction before execution.

7. A computerized method, comprising:

5 obtaining a memory request;
storing the memory request in a first memory request container; and
processing the memory request from the first memory request container by a
cache controller, wherein the processing includes
identifying a type of the memory request,
10 processing the memory request in a local cache as a function of a first
condition, and
processing the memory request using a second memory request
container as a function of a second condition,
wherein the processing of the memory request using the second
15 memory request container includes updating a portion of the local cache with a
portion of a remote cache.

8. The computerized method of claim 7, wherein the obtaining of the memory request includes obtaining a memory load or a memory store request.

20 9. The computerized method of claim 7, wherein the storing of the memory request in the first memory request container includes storing the memory request in an Initial Request Queue (IRQ).

25 10. The computerized method of claim 7, wherein the processing of the memory request in the local cache as a function of the first condition includes processing the memory request in the local cache when the local cache includes an entry associated with the memory request.

11. The computerized method of claim 7, wherein the processing of the memory request using a second memory request container as a function of a second condition includes processing the memory request using a Forced Order Queue (FOQ) as a function of the second condition.

5

12. The computerized method of claim 11, wherein the processing of the memory request using the FOQ as a function of the second condition includes processing the memory request using the FOQ when the memory request matches a corresponding request in the FOQ.

10

13. The computerized method of claim 11, wherein the processing of the memory request using the FOQ as a function of the second condition includes adding the memory request to the FOQ when the memory request is not initially processed by the local cache, and when the memory request does not initially match 15 a corresponding request already in the FOQ.

15

14. The computerized method of claim 11, wherein the processing of the memory request using the FOQ as a function of the second condition includes processing the memory request in the FOQ when local cache processing is bypassed.

20

15. The computerized method of claim 11, wherein the processing of the memory request using the FOQ as a function of the second condition includes processing the memory request in the FOQ when the memory request includes a synchronization request that causes local cache processing to be bypassed.

25

16. A computerized method, comprising:
adding a set of memory instructions of a first type to a memory instruction container;

executing the set of memory instructions of the first type from the memory instruction container;

receiving a memory instruction of a second type;

receiving a synchronization instruction that temporarily blocks further

5 execution of memory instructions of the first type from the memory instruction container;

adding an additional set of memory instructions of the first type to the memory instruction container;

executing the memory instruction of the second type; and

10 upon such execution of the memory instruction of the second type, executing the additional set of memory instructions of the first type from the memory instruction container.

17. The computerized method of claim 16, wherein the adding of the set of

15 memory instructions of the first type to the memory instruction container includes adding the set of memory instructions of the first type to a memory instruction queue.

18. The computerized method of claim 16, wherein the adding of the set of

20 memory instructions of the first type to the memory instruction container includes adding a set of scalar memory instructions to the memory instruction container.

19. The computerized method of claim 16, wherein the executing of the set of

memory instructions of the first type from the memory instruction container includes

25 accessing a local or remote cache unit.

20. The computerized method of claim 16, wherein the receiving of the memory instruction of the second type includes receiving a vector memory instruction.

21. The computerized method of claim 16, wherein the executing of the memory instruction of the second type includes accessing a remote cache unit, and updating a local cache unit with contents from the remote cache unit.

5 22. A scalar processor, comprising:

a cache; and

a scalar load/store unit having a first container, a second container, and a cache controller,

wherein the cache controller obtains a scalar load/store command from the 10 first container, the scalar load/store command having a scalar load/store instruction and one or more addresses,

wherein the cache services the scalar load/store command when the scalar load/store command hits in the cache,

wherein the scalar load/store command is added to the second container 15 when the scalar load/store command misses in the cache, and wherein one or more lines in the cache are allocated by a remote cache when the scalar load/store command is added to the second container.

23. The scalar processor of claim 22, wherein the first container comprises an 20 Initial Request Queue (IRQ), and wherein the second container comprises a Forced Order Queue (FOQ).

24. The scalar processor of claim 22, wherein the scalar load/store unit further 25 includes an address generator to generate one or more physical addresses from the one or more addresses of the scalar load/store command.

25. The scalar processor of claim 24, wherein the address generator generates the one or more physical addresses using a translation look-aside buffer (TLB).

26. A processing unit, comprising:
a remote cache interface; and
a scalar processor having a local cache,
wherein the scalar processor dispatches a memory instruction, obtains one or
5 more address operands, creates a virtual address from the one or more address
operands, translates the virtual address into a physical address using a translation
look-aside buffer (TLB), and executes the memory instruction using a cache
controller, the cache controller determining whether to obtain data from the local
cache or to allocate one or more lines of the local cache through the remote cache
10 interface.

27. A processing unit, comprising:
a cache interface;
a vector processor; and
15 a scalar processor,
wherein the scalar processor processes a plurality of scalar memory
instructions in a memory instruction container,
wherein the vector processor receives a vector memory instruction,
wherein the scalar processor receives a synchronization instruction and
20 temporarily blocks further processing of scalar memory instructions in the memory
instruction container,
wherein the vector processor processes the vector memory instruction by
accessing a cache through the cache interface, and
wherein the scalar processor continues to process further scalar memory
25 instructions in the memory instruction container after the vector processor has
processed the vector memory instruction.

28. The processing unit of claim 27, wherein the memory instruction container
comprises a memory instruction queue.